

CLAIMS

What is claimed is:

1. A processor-based method for allocating resources to a plurality of applications, comprising:
 - determining available resources of a networked computing system;
 - determining, for each application, required resources of the application;
 - determining an assigned subset of the available resources for each application as a function of the required resources of the application and the available resources, wherein the function reduces communication delays between resources of the subset of the available resources in conformance with bandwidth capacity requirements of the application and in conformance with network bandwidth limitations; and
 - associating the applications with the assigned subsets of resources.
2. The method of claim 1, wherein the available resources comprise processing resources, networking resources, and storage resources.
3. The method of claim 2, wherein the processing resources comprise servers each having at least one processor.
4. The method of claim 2, wherein the networking resources comprise network switches.
5. The method of claim 2, wherein the storage resources comprise a storage area network.
6. The method of claim 5, wherein the storage area network includes at least one pair of redundant core switches coupled to storage devices, the core switches coupled to the processing resources via a plurality of edge switches.
7. The method of claim 1, wherein reducing the communications delay between resources comprises solving a mixed-integer programming problem.

200313904-1 (HPCO.146PA)

8. The method of claim 7, wherein the available resources include network switches coupled with the available resources, and the mixed-integer programming problem reduces communication delays between resources of the subset of the available resources by reducing data traffic on network links that interconnect the network switches.

9. A system comprising:

means for determining available resources of a networked computing system;

means for determining required resources for each application of a plurality of applications;

means for determining an assigned subset of the available resources for each application as a function of the required resources of the application and the available resources, wherein the function reduces communication delays between resources of the subset of the available resources in conformance with bandwidth capacity requirements of the application and in conformance with network bandwidth limitations; and;

means for associating the applications with the assigned subsets of resources.

10. A computer-readable medium configured with instructions for causing a processor of a data processing arrangement to allocate resources to a plurality of applications, comprising:

determining available resources of a networked computing system;

determining, for each application, required resources of the application;

determining an assigned subset of the available resources for each application as a function of the required resources of the application and the available resources, wherein the function reduces communication delays between resources of the subset of the available resources in conformance with bandwidth capacity requirements of the application and in conformance with network bandwidth limitations; and

associating the applications with the assigned subsets of resources.

11. The computer-readable medium of claim 10, wherein the available resources comprise processing resources, networking resources, and storage resources.

200313904-1 (HPCO.146PA)

12. The computer-readable medium of claim 11, wherein the processing resources comprise servers each having at least one processor.

13. The computer-readable medium of claim 11, wherein the networking resources comprise network switches.

14. The computer-readable medium of claim 11, wherein the storage resources comprise a storage area network.

15. The computer-readable medium of claim 14, wherein the storage area network includes at least one pair of redundant core switches coupled to storage devices, the core switches coupled to the processing resources via a plurality of edge switches.

16. The computer-readable medium of claim 10, wherein reducing the communications delay between resources comprises solving a mixed-integer programming problem.

17. The computer-readable medium of claim 16, wherein the available resources include network switches coupled with the processing resources, and the mixed-integer programming problem reduces communication delays between resources of the subset of the available resources by reducing data traffic on network links that interconnect the network switches.

18. A system, comprising:
 - a plurality of network-coupled processing resources;
 - a plurality of storage resources network-coupled to the processing resources,wherein the processing and storage resources are allocated to a plurality of applications;
 - a computing arrangement configured to,
 - determine, for each application of the plurality of applications, required resources of the application;
 - determining an assigned subset of the processing and storage resources for each application as a function of the required resources of the application and the processing and storage resources, wherein the function reduces communication delays between resources of the subset of the network and processing resources in conformance with bandwidth capacity requirements of the application and in conformance with network bandwidth limitations;
 - associate the applications with the assigned subsets of processing and storage resources.
19. The system of claim 18, wherein the processing resources comprise servers each having at least one processor.
20. The system of claim 18, wherein the storage resources comprise a storage area network.
21. The system of claim 20, wherein the storage area network includes at least one pair of redundant core switches coupled to storage devices, the core switches coupled to the network via a plurality of edge switches.
22. The system of claim 18, wherein the computing arrangement is configured to reduce the communications delay between resources by solving a mixed-integer programming problem.

200313904-1 (HPCO.146PA)

23. The system of claim 22, wherein processing resources are coupled by network switches, and the mixed-integer programming problem reduces communication delays between resources by reducing data traffic on network links that interconnect the network switches.